

Date of Deposit: April 9, 2001

I hereby certify that this paper is being deposited with the United States Postal Service "EXPRESS MAIL Post Office to Addressee" service under 37 CFR 1.10 on the date indicated above and is addressed to the Assistant Commissioner for Patents; Washington, DC 20231.

Kelli J. Withrow

[illegible]

Kelly [Signature]
Signature

APPLICATION FOR UNITED STATES LETTERS PATENT

Title: METHOD FOR DETERMINING THE SHAPE OF OBJECTS
DIRECTLY FROM RANGE IMAGES

Inventors: Sarah F. Frisken
Ronald N. Perry

[illegible]

Method for Determining the Shape of Objects Directly From Range Images

Field of the Invention

5

The present invention relates generally to the field of computer graphics, and particularly to generating models of graphics objects from range images.

Background of the Invention

10

Several techniques have been developed in research and industry for determining 3D geometry from a set of aligned range images, see Curlass "*From range scans to 3D models*," Comp. Graphics, Vol.33, No.4, 1999.

09829430-010904
TOP SECRET

15

Recent work has shown that volumetric methods are robust to scanner noise and alignment uncertainty and provide good quality, water-tight models, see Curlass et al. "*A volumetric method for building complex models from range images*," Proc. SIGGRAPH'96, pp. 303-312, 1996, Hilton et al. "*Reliable surface reconstruction from multiple range images*," 4th Euro. Conf. Comp. Vis., pp. 117-126, 1996, and
20 Wheeler "*Automatic modeling and localization for object recognition*", Ph.D. Thesis, Carnegie Mellon University, 1996.

25

However, most of these methods suffer from limited resolution, large memory requirements and long processing time. In addition, these methods produce mesh models with an excessively large number of triangle.

These methods construct a range surface for each aligned range image, and fill a fixed resolution volumetric representation with signed distances from the range surfaces. The methods use various techniques to reduce the time required to fill and access the volumetric data, including run length encoding of the distance values, binary encoding of regions outside a bounded region of the surface, and a 3-color octree representation of the volume. The distance values from multiple scans can be combined probabilistically using order-independent, or incremental updating. These methods all build a triangle model of the iso-surface of the distance volume using a well known Marching Cubes technique, see Lorensen et al. in “*Marching Cubes: A High Resolution 3D Surface Construction Algorithm*,” Proc. SIGGRAPH '87, pp.163-169, 1987.

Constructing 3D range surfaces and computing distances from these iso-surfaces significantly increases computational requirements. If, instead, the distance field could be generated directly from 2D range images, model generation times could be greatly reduced.

However, range images do not provide true distance values. In the simplest case, a range image records the perpendicular projected distance from the object's surface to the image plane. The projected distance field is the same as the true distance field in only two special cases: throughout the distance field for a planar surface exactly parallel to the image plane, and at the surface where both true distances and the projected distances are zero for any surface.

Other than the first case, the projected distance field differs from the true distance field for points off the surface. This results in artifacts when projected distance fields from different viewpoints are combined.

Prior art described in U.S. Patent Application Sn. 09/206,071 "*A Method for Estimating Volumetric Distance Maps from 2D Depth Images*" filed by Frisker et al. on December 4, 1999 showed that a 3D projected distance field can be

corrected by dividing sampled distances by the magnitude of the local 3D gradient of the projected distance field. This results in a better approximation of the true distance field near the surface, yielding better results when combining projected distance fields. However, computing the 3D gradient of the projected distance field directly from sampled distance values can be prohibitively expensive since it can require six or more additional evaluations of the distance field per gradient.

Therefore, it is desired to determine corrective distance field gradients directly from range images rather than from 3D volumetric data.

Summary of the Invention

The present invention provides a method for determining a 2D gradient magnitude image from a range image of an object from which the 3D gradient magnitude can be derived. The range image includes a plurality of intensity values at pixel locations. The intensity values correspond to distances to a surface of the object.

The method scales the range image so that a unit intensity value corresponds to a unit distance value. Then, for each pixel $s_{i,j}$ in the range image, the method determines a horizontal central difference dx , and a vertical central difference dy . Then, the 2D gradient magnitude at each pixel $s_{i,j}$ is set to one half of the square root of $(dx^2 + dy^2 + 4)$.

The magnitude of the gradient of a 3D projected distance field can then be determined at a 3D point \mathbf{p} by perpendicularly projecting the point \mathbf{p} onto a location (x, y) of the scaled range image, by interpolating the gradient magnitude at (x, y) from the corresponding 2D gradient magnitude image values near the location (x, y) , and by setting the 3D gradient magnitude at \mathbf{p} to the interpolated 2D gradient magnitude at (x, y) .

Brief Description of the Drawings

Figure 1 is a block diagram of a method for determining a 2D gradient magnitude image according to the invention; and

Figure 2 is a block diagram of a method for determining the magnitude of the gradient of a 3D projected distance field.

Detailed Description of the Preferred Embodiment

Introduction

We describe a method for determining the magnitude of the 3D gradient of the projected distance field directly from 2D range images. Using our method, the 3D shape of an object can be derived directly from 2D range images rather than range surfaces as in the prior art. Furthermore, this method can be used to generate an adaptively sampled distance field (ADF) rather than a distance volume or a three-color octree, resulting in a significant savings in memory and distance computations.

We correct the 3D projected distance field by dividing sampled distances by the local gradient magnitude. This results in a better approximation of the true distance field near the surface, yielding better results when combining projected distance fields. We derive the 3D gradient from a 2D gradient image generated during a pre-processing step, resulting in significantly faster generation.

Because range images of some scanning systems are not simple projected distances, the range images must first be converted to this form. While this conversion results in some loss of information from the scan, the benefits include significantly faster performance, thus allowing interactive updating during data acquisition.

Adaptively Sampled Distance Fields

We prefer to use an adaptively sampled distance field (ADF) to represent the shape of an object. Here, the shape of objects is acquired by range images 105, see Figure 1. The basic data structure of an ADF is described in U.S. Patent Application Sn. 09/370,091 "*Detail-Directed Distance Fields*" filed by Frisken et al. on August 6, 1999, incorporated herein in its entirety by reference.

ADFs adaptively sample a signed distance field of an object and store the sample values in a spatial hierarchy for fast processing. Distance values with a positive sign are exterior to the object, negative distance values are interior, and zero distance values represent the surface of the object.

Methods that can operate on ADFs are described in: U.S. Patent Applications Sn. 09/810,983 "*System and Method for Generating Adaptively Sampled Distance*

Fields with Bounded Distance Trees” filed by Perry et al. on March 16, 2001, U.S. Patent Application Sn. 09/810,839 “*Conversion of Adaptively Sampled Distance Fields to Triangles*” filed by Frisken et al. on March 16, 2001, U.S. Patent Application Sn. 09/811,010 “*System and Method for Modeling Graphics Objects*”
5 filed by Perry et. al. on March 16, 2001, and U.S. Patent Application 09/809,682 “*System and Method for Converting Range Data to 3D Models,*” filed by Frisken et al. on March 16, 2001.

Frisken et al., in “*Adaptively sampled distance fields: a general representation of shape for computer graphics,*” Proc. SIGGRAPH'00, pp. 249-254, 2000, found that
10 even in 2D, ADFs require twenty times fewer distance computations than a comparable three color octree representation such as is used by Wheeler, in “*Automatic modeling and localization for object recognition*”, Ph.D. Thesis, Carnegie Mellon University, 1996. In addition, ADFs can be interactively edited
15 via a sculpting interface so that holes and other surface anomalies from occlusions and sensor noise can be easily corrected, see Perry et al. “*Kizamu: A system for sculpting digital characters,*” Proc. SIGGRAPH'01, 2001, and U.S. Patent Application 09/810,261 “*System and Method for Sculpting Digital Models*” filed
20 by Perry et al. on March 16, 2001.

2D Gradient Magnitude Images from Range Images

Figure 1 shows a method for determining a 2D gradient magnitude image I_{GM} 125
from a range image I_R 105. In step 110, the intensity value at each pixel of the
25 range image 105 is scaled to produce a scaled range image I_{RS} 115. The scaling ensures that a normalized unit intensity value at each pixel represents a unit distance value.

Step 120 determines the gradient magnitude image 125 from the scaled range image 115.

- 5 For each pixel (i,j), we determine central difference values:

$$dx = I_{RS}(i+1, j) - I_{RS}(i-1, j)$$

$$dy = I_{RS}(i, j+1) - I_{RS}(i, j-1)$$

and the 2D gradient magnitude image value at (i,j):

10 $I_{GM}(i,j) = 0.5 * \text{sqrt}(dx^2 + dy^2 + 4).$

Magnitude of Gradients in a 3D Projected Distance Field

Figure 2 shows a method for determining a magnitude of a gradient 140 at a 3D point **p** 130 from the scaled range image 115 and the gradient magnitude image 125. First, a perpendicular projection (x,y) of **p** 130 onto the scaled range image 115 is determined. Next, a gradient magnitude at (x,y) is interpolated from the corresponding values of the 2D gradient magnitude image 125 near the location (x,y). Any known prior art interpolation method, such as bilinear interpolation, can
20 be used. Finally, the magnitude of the gradient 140 at **p** 130 is set to the interpolated gradient magnitude at (x,y).

This invention is described using specific terms and examples. It is to be understood that various other adaptations and modifications may be made within
25 the spirit and scope of the invention. Therefore, it is the object of the appended claims to cover all such variations and modifications as come within the true spirit and scope of the invention.